

1226 4/20/06

## Refine Search

---

### Search Results -

Terms	Documents
L9 and (auxiliar\$ same (partition\$ or cluster\$3) same node)	1

---

**Database:** US Pre-Grant Publication Full-Text Database  
 US Patents Full-Text Database  
 US OCR Full-Text Database  
 EPO Abstracts Database  
 JPO Abstracts Database  
 Derwent World Patents Index  
 IBM Technical Disclosure Bulletins

**Search:** L10

---

### Search History

---

**DATE:** Thursday, April 20, 2006 [Printable Copy](#) [Create Case](#)

**Set Name** **Query**

side by side

**Hit Count** **Set Name**  
result set

<u>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</u>		
<u>L10</u>	L9 and (auxiliar\$ same (partition\$ or cluster\$3) same node)	1 <u>L10</u>
<u>L9</u>	L7 and (base near relation)	12 <u>L9</u>
<u>L8</u>	L7 not L6	0 <u>L8</u>
<u>L7</u>	(L1 or L3) and L4	40 <u>L7</u>
<u>L6</u>	(L1 or L3) and L5	40 <u>L6</u>
<u>L5</u>	tuple\$1 same (table\$1 or row\$1 or column\$1) same join\$3	634 <u>L5</u>
<u>L4</u>	tuple\$1 same (table\$1 or row\$1 or column\$1) same join\$	634 <u>L4</u>
<u>L3</u>	L2 not L1	28 <u>L3</u>
<u>L2</u>	(materializ\$3 or materialis\$3) near view\$1	374 <u>L2</u>
<u>L1</u>	"materialized view"	346 <u>L1</u>

END OF SEARCH HISTORY

**Refine Search****Search Results -**

Terms	Documents
L11 and (join near attribute)	2

**Database:** US Pre-Grant Publication Full-Text Database  
 US Patents Full-Text Database  
 US OCR Full-Text Database  
 EPO Abstracts Database  
 JPO Abstracts Database  
 Derwent World Patents Index  
 IBM Technical Disclosure Bulletins

**Search:** L12

**Search History**

**DATE:** Thursday, April 20, 2006 [Printable Copy](#) [Create Case](#)

<u>Set</u> <u>Name</u>	<u>Query</u>	<u>Hit</u> <u>Count</u>	<u>Set</u> <u>Name</u>
side by side			result set
DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
<u>L12</u>	L11 and (join near attribute)	2	<u>L12</u>
<u>L11</u>	(L1 or L3) and (auxiliar\$ same (partition\$ or cluster\$3 or group\$3 or categoriz\$3 or categoris\$3))	5	<u>L11</u>
<u>L10</u>	L9 and (auxiliar\$ same (partition\$ or cluster\$3) same node)	1	<u>L10</u>
<u>L9</u>	L7 and (base near relation)	12	<u>L9</u>
<u>L8</u>	L7 not L6	0	<u>L8</u>
<u>L7</u>	(L1 or L3) and L4	40	<u>L7</u>
<u>L6</u>	(L1 or L3) and L5	40	<u>L6</u>
<u>L5</u>	tuple\$1 same (table\$1 or row\$1 or column\$1) same join\$3	634	<u>L5</u>
<u>L4</u>	tuple\$1 same (table\$1 or row\$1 or column\$1) same join\$	634	<u>L4</u>
<u>L3</u>	L2 not L1	28	<u>L3</u>
<u>L2</u>	(materializ\$3 or materialis\$3) near view\$1	374	<u>L2</u>
<u>L1</u>	"materialized view"	346	<u>L1</u>

END OF SEARCH HISTORY

# Refine Search

---

## Search Results -

Terms	Documents
L19 and (join near attribute)	1

---

**Database:**

- US Pre-Grant Publication Full-Text Database
- US Patents Full-Text Database
- US OCR Full-Text Database
- EPO Abstracts Database
- JPO Abstracts Database
- Derwent World Patents Index
- IBM Technical Disclosure Bulletins

**Search:**

---

## Search History

---

**DATE:** Thursday, April 20, 2006    [Printable Copy](#)    [Create Case](#)

<u>Set</u>	<u>Name</u>	<u>Query</u>	<u>Hit Count</u>	<u>Set</u>
side by side				result set
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>				
<u>L24</u>	L19 and (join near attribute)		1	<u>L24</u>
<u>L23</u>	L20 and (join near attribute)		2	<u>L23</u>
<u>L22</u>	L21 and (join near attribute)		6	<u>L22</u>
<u>L21</u>	L16 and (base near relation)		27	<u>L21</u>
<u>L20</u>	L17 and (base near relation)		10	<u>L20</u>
<u>L19</u>	L18 and (base near relation)		7	<u>L19</u>
<u>L18</u>	L15 and l2		47	<u>L18</u>
<u>L17</u>	L14 and l2		108	<u>L17</u>
<u>L16</u>	L13 and l2		214	<u>L16</u>
<u>L15</u>	707/200-206.ccls.		6980	<u>L15</u>
<u>L14</u>	707/100-104.1.ccls.		13609	<u>L14</u>
<u>L13</u>	707/1-10.ccls.		20502	<u>L13</u>
<u>L12</u>	L11 and (join near attribute)		2	<u>L12</u>

<u>L11</u>	(L1 or L3) and (auxiliar\$ same (partition\$ or cluster\$3 or group\$3 or categoriz\$3 or categoris\$3))	5	<u>L11</u>
<u>L10</u>	L9 and (auxiliar\$ same (partition\$ or cluster\$3) same node)	1	<u>L10</u>
<u>L9</u>	L7 and (base near relation)	12	<u>L9</u>
<u>L8</u>	L7 not L6	0	<u>L8</u>
<u>L7</u>	(L1 or L3) and L4	40	<u>L7</u>
<u>L6</u>	(L1 or L3) and L5	40	<u>L6</u>
<u>L5</u>	tuple\$1 same (table\$1 or row\$1 or column\$1) same join\$3	634	<u>L5</u>
<u>L4</u>	tuple\$1 same (table\$1 or row\$1 or column\$1) same join\$	634	<u>L4</u>
<u>L3</u>	L2 not L1	28	<u>L3</u>
<u>L2</u>	(materializ\$3 or materialis\$3) near view\$1	374	<u>L2</u>
<u>L1</u>	"materialized view"	346	<u>L1</u>

END OF SEARCH HISTORY

120410/6

## Refine Search

---

### Search Results -

Terms	Documents
L5 and (parallel near (database or (data near base)))	2

---

**Database:**

- US Pre-Grant Publication Full-Text Database
- US Patents Full-Text Database
- US OCR Full-Text Database
- EPO Abstracts Database
- JPO Abstracts Database
- Derwent World Patents Index
- IBM Technical Disclosure Bulletins

**Search:**

Refine Search

Recall Text
Clear
Interrupt

---

### Search History

---

**DATE:** Thursday, April 20, 2006    [Printable Copy](#)    [Create Case](#)

**Set Name** **Query**  
side by side

**Hit Count** **Set Name**  
result set

<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>		
<u>L6</u>	L5 and (parallel near (database or (data near base)))	2 <u>L6</u>
<u>L5</u>	L4 and L3	7 <u>L5</u>
<u>L4</u>	L1 and L2	37 <u>L4</u>
<u>L3</u>	join near attribute	180 <u>L3</u>
<u>L2</u>	base near relation	18221 <u>L2</u>
<u>L1</u>	(materializ\$3 or materialis\$3) near view\$1	374 <u>L1</u>

END OF SEARCH HISTORY

## Hit List

---

<a href="#">First Hit</a>	<a href="#">Clear</a>	<a href="#">Generate Collection</a>	<a href="#">Print</a>	<a href="#">Fwd Refs</a>	<a href="#">Bkwd Refs</a>
<a href="#">Generate OACS</a>					

**Search Results - Record(s) 1 through 2 of 2 returned.**

---

1. Document ID: US 6990503 B1

**Using default format because multiple data bases are involved.**

L6: Entry 1 of 2

File: USPT

Jan 24, 2006

US-PAT-NO: 6990503

DOCUMENT-IDENTIFIER: US 6990503 B1

TITLE: Rescheduling transactions in a database system

DATE-ISSUED: January 24, 2006

**INVENTOR-INFORMATION:**

NAME	CITY	STATE	ZIP CODE	COUNTRY
Luo; Gang	Madison	WI		US
Ellmann; Curt J.	Madison	WI		US
Naughton; Jeffrey F.	Madison	WI		US
Watzke; Michael W.	Madison	WI		US

US-CL-CURRENT: 707/200

<a href="#">Full</a>	<a href="#">Title</a>	<a href="#">Citation</a>	<a href="#">Front</a>	<a href="#">Review</a>	<a href="#">Classification</a>	<a href="#">Date</a>	<a href="#">Reference</a>	<a href="#">Sequences</a>	<a href="#">Attachments</a>	<a href="#">Claims</a>	<a href="#">KINIC</a>	<a href="#">Drawn D</a>
----------------------	-----------------------	--------------------------	-----------------------	------------------------	--------------------------------	----------------------	---------------------------	---------------------------	-----------------------------	------------------------	-----------------------	-------------------------

- 
2. Document ID: US 6952692 B1

L6: Entry 2 of 2

File: USPT

Oct 4, 2005

US-PAT-NO: 6952692

DOCUMENT-IDENTIFIER: US 6952692 B1

TITLE: Execution of requests in a parallel database system

<a href="#">Full</a>	<a href="#">Title</a>	<a href="#">Citation</a>	<a href="#">Front</a>	<a href="#">Review</a>	<a href="#">Classification</a>	<a href="#">Date</a>	<a href="#">Reference</a>	<a href="#">Sequences</a>	<a href="#">Attachments</a>	<a href="#">Claims</a>	<a href="#">KINIC</a>	<a href="#">Drawn D</a>
----------------------	-----------------------	--------------------------	-----------------------	------------------------	--------------------------------	----------------------	---------------------------	---------------------------	-----------------------------	------------------------	-----------------------	-------------------------

<a href="#">Clear</a>	<a href="#">Generate Collection</a>	<a href="#">Print</a>	<a href="#">Fwd Refs</a>	<a href="#">Bkwd Refs</a>	<a href="#">Generate OACS</a>
-----------------------	-------------------------------------	-----------------------	--------------------------	---------------------------	-------------------------------

Terms	Documents
L5 and (parallel near (database or (data near base)) )	2

**Display Format:** [- ]

[Previous Page](#)    [Next Page](#)    [Go to Doc#](#)

# Refine Search

---

## Search Results -

Terms	Documents
L12 and L1	1

---

**Database:**

- US Pre-Grant Publication Full-Text Database
- US Patents Full-Text Database
- US OCR Full-Text Database
- EPO Abstracts Database
- JPO Abstracts Database
- Derwent World Patents Index
- IBM Technical Disclosure Bulletins

**Search:**

▲
▼

Refine Search

  

Recall Text
Clear
Interrupt

---

## Search History

---

**DATE:** Thursday, April 20, 2006    [Printable Copy](#)    [Create Case](#)

<u>Set Name</u> <u>Query</u>	<u>Hit Count</u>	<u>Set Name</u>
result set		
<u>L13</u> L12 and L1	1	<u>L13</u>
<u>L12</u> L11 and ("join attribute")	6	<u>L12</u>
<u>L11</u> L10 and L9	11	<u>L11</u>
<u>L10</u> "base relation"	3115	<u>L10</u>
<u>L9</u> "parallel database"	702	<u>L9</u>
<u>L8</u> L7 and L6	2	<u>L8</u>
<u>L7</u> 707/\$.ccls.	33933	<u>L7</u>
<u>L6</u> L5 and (parallel near (database or (data near base)))	2	<u>L6</u>
<u>L5</u> L4 and L3	7	<u>L5</u>
<u>L4</u> L1 and L2	37	<u>L4</u>
<u>L3</u> join near attribute	180	<u>L3</u>
<u>L2</u> base near relation	18221	<u>L2</u>
<u>L1</u> (materializ\$3 or materialis\$3) near view\$1	374	<u>L1</u>

END OF SEARCH HISTORY

## Hit List

First Hit	Clear	Generate Collection	Print	Fwd Refs	Bkwd Refs
		Generate OACS			

Search Results - Record(s) 1 through 1 of 1 returned.

1. Document ID: US 6990503 B1

Using default format because multiple data bases are involved.

L13: Entry 1 of 1

File: USPT

Jan 24, 2006

US-PAT-NO: 6990503

DOCUMENT-IDENTIFIER: US 6990503 B1

TITLE: Rescheduling transactions in a database system

DATE-ISSUED: January 24, 2006

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Luo; Gang	Madison	WI		US
Ellmann; Curt J.	Madison	WI		US
Naughton; Jeffrey F.	Madison	WI		US
Watzke; Michael W.	Madison	WI		US

US-CL-CURRENT: 707/200

Full	Title	Citation	Front	Review	Classification	Date	Reference	Sequences	Attachments	Claims	KWIC	Drawn D.
------	-------	----------	-------	--------	----------------	------	-----------	-----------	-------------	--------	------	----------

Clear	Generate Collection	Print	Fwd Refs	Bkwd Refs	Generate OACS
-------	---------------------	-------	----------	-----------	---------------

Terms	Documents
L12 and L1	1

Display Format: [-] Change Format

[Previous Page](#)

[Next Page](#)

[Go to Doc#](#)

SEARCH 4/20/06

## Refine Search

---

### Search Results -

Terms	Documents
L13 and node	2

---

**Database:** US Pre-Grant Publication Full-Text Database  
 US Patents Full-Text Database  
 US OCR Full-Text Database  
 EPO Abstracts Database  
 JPO Abstracts Database  
 Derwent World Patents Index  
 IBM Technical Disclosure Bulletins

**Search:** L14

Refine Search
Recall Text
Clear
Interrupt

---

### Search History

---

**DATE:** Thursday, April 20, 2006 [Printable Copy](#) [Create Case](#)

**Set Name Query**

**Hit Count Set Name**  
result set

side by side

<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>	
<u>L14</u> L13 and node	2 <u>L14</u>
<u>L13</u> L12 and (parallel near database)	3 <u>L13</u>
<u>L12</u> L11 and L2	35 <u>L12</u>
<u>L11</u> 707/\$.ccls.	33933 <u>L11</u>
<u>L10</u> L8 and (parallel near database)	1 <u>L10</u>
<u>L9</u> L8 and (receiv\$3 same tuple) and (parallel near database)	0 <u>L9</u>
<u>L8</u> L6 and tuple	6 <u>L8</u>

*DB=USPT; PLUR=YES; OP=OR*

<u>L7</u> L6	5 <u>L7</u>
--------------	-------------

*DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR*

<u>L6</u> L4 and ("join attribute")	6 <u>L6</u>
-------------------------------------	-------------

*DB=USPT; PLUR=YES; OP=OR*

<u>L5</u> L4	15 <u>L5</u>
--------------	--------------

*DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR*

<u>L4</u>	L3 and node	23	<u>L4</u>
<u>L3</u>	L2 and join	32	<u>L3</u>
<u>L2</u>	L1 and ("base relation")	37	<u>L2</u>
<u>L1</u>	(materializ\$3 or materialis\$3) near view	374	<u>L1</u>

END OF SEARCH HISTORY

# Refine Search

---

## Search Results -

Terms	Documents
L8 and (parallel near database)	1

---

**Database:**

- US Pre-Grant Publication Full-Text Database
- US Patents Full-Text Database
- US OCR Full-Text Database
- EPO Abstracts Database
- JPO Abstracts Database
- Denwent World Patents Index
- IBM Technical Disclosure Bulletins

**Search:**

[]
[]
Refine Search

Recall Text
Clear
Interrupt

---

## Search History

---

**DATE:** Thursday, April 20, 2006    [Printable Copy](#)    [Create Case](#)

**Set Name** **Query**

side by side

**Hit Count** **Set Name**

result set

*DB=PGPB, USPT, USOC, EPAB, JPAB, DWPI, TDBD; PLUR=YES; OP=OR*

L10	L8 and (parallel near database)	1	<u>L10</u>
-----	---------------------------------	---	------------

L9	L8 and (receiv\$3 same tuple) and (parallel near database)	0	<u>L9</u>
----	--	---	-----------

L8	L6 and tuple	6	<u>L8</u>
----	--------------	---	-----------

*DB=USPT; PLUR=YES; OP=OR*

L7	L6	5	<u>L7</u>
----	----	---	-----------

*DB=PGPB, USPT, USOC, EPAB, JPAB, DWPI, TDBD; PLUR=YES; OP=OR*

L6	L4 and ("join attribute")	6	<u>L6</u>
----	---------------------------	---	-----------

*DB=USPT; PLUR=YES; OP=OR*

L5	L4	15	<u>L5</u>
----	----	----	-----------

*DB=PGPB, USPT, USOC, EPAB, JPAB, DWPI, TDBD; PLUR=YES; OP=OR*

L4	L3 and node	23	<u>L4</u>
----	-------------	----	-----------

L3	L2 and join	32	<u>L3</u>
----	-------------	----	-----------

L2	L1 and ("base relation")	37	<u>L2</u>
----	--------------------------	----	-----------

L1	(materializ\$3 or materialis\$3) near view	374	<u>L1</u>
----	--	-----	-----------

END OF SEARCH HISTORY

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

**Search Results****BROWSE****SEARCH****IEEE XPLORE GUIDE**

Results for "(((materializ\* <paragraph> view <paragraph> tuple) <and> cluster\*)<in>metadata..."  
Your search matched 0 documents.

 [e-mail](#)

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

**» Search Options**[View Session History](#)**Modify Search**[New Search](#) Check to search only within this results setDisplay Format:  Citation  Citation & Abstract**» Key****IEEE JNL** IEEE Journal or Magazine**IEE JNL** IEE Journal or Magazine**IEEE CNF** IEEE Conference Proceeding**IEE CNF** IEE Conference Proceeding**IEEE STD** IEEE Standard**No results were found.**

Please edit your search criteria and try again. Refer to the Help pages if you need assistance.  
search.

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by  
 Inspec®



Home | Login | Logout | Access Information | Alerts |

Welcome United States Patent and Trademark Office

### Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "((materializ\* <paragraph> view <paragraph> tuple)<in>metadata)"

e-mail

Your search matched 0 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

#### » Search Options

[View Session History](#)

[New Search](#)

#### Modify Search

((materializ\* <paragraph> view <paragraph> tuple)<in>metadata)

Search

Check to search only within this results set

Display Format:  Citation  Citation & Abstract

#### » Key

**IEEE JNL** IEEE Journal or Magazine

**IEE JNL** IEE Journal or Magazine

**IEEE CNF** IEEE Conference Proceeding

**IEE CNF** IEE Conference Proceeding

**IEEE STD** IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance.

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by  
 Inspec


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

**Search Results****BROWSE****SEARCH****IEEE XPLORE GUIDE**

Results for "(((materializ\* &lt;near&gt; view) &lt;and&gt; (base &lt;near&gt; relation) &lt;and&gt; (parallel &lt;...&gt;"

[e-mail](#)

Your search matched 3 of 1340257 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending** order.**» Search Options**[View Session History](#)[New Search](#)**Modify Search**

[Search](#) Check to search only within this results setDisplay Format:  Citation  Citation & Abstract
 [Select All](#) [Deselect All](#)
 1. **A comparison of three methods for join view maintenance in parallel RDE**

Luo, G.; Naughton, J.F.; Ellmann, C.J.; Watzke, M.W.;  
Data Engineering, 2003. Proceedings. 19th International Conference on  
 5-8 March 2003 Page(s):177 - 188

[AbstractPlus](#) | Full Text: [PDF\(634 KB\)](#) IEEE CNF  
[Rights and Permissions](#)
 2. **Parallel generation of base relation snapshots for materialized view main warehouse environment**

Saeki, S.; Bhalla, S.; Hasegawa, M.;  
Parallel Processing Workshops, 2002. Proceedings. International Conference  
 18-21 Aug. 2002 Page(s):383 - 390  
 Digital Object Identifier 10.1109/ICPPW.2002.1039755

[AbstractPlus](#) | Full Text: [PDF\(285 KB\)](#) IEEE CNF  
[Rights and Permissions](#)
 3. **Making views self-maintainable for data warehousing**

Quass, D.; Gupta, A.; Mumick, I.S.; Widom, J.;  
Parallel and Distributed Information Systems, 1996., Fourth International Conference  
 18-20 Dec. 1996 Page(s):158 - 169  
 Digital Object Identifier 10.1109/PDIS.1996.568677

[AbstractPlus](#) | Full Text: [PDF\(1228 KB\)](#) IEEE CNF  
[Rights and Permissions](#)
[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –

**Indexed by**



[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

## Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "(((materializ\* <near> view) <and> (base <near> relation) <and> (parallel <...>))

[e-mail](#)

Your search matched 0 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

### » Search Options

[View Session History](#)

### Modify Search

[New Search](#)

Check to search only within this results set

Display Format:  Citation  Citation & Abstract

### » Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance.

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by  
 Inspec

/next page

 **PORTAL**  
 USPTO

Subscribe (Full Service) [Register \(Limited Service, Free\)](#) [Login](#)
Search:  The ACM Digital Library  The Guide

"receiving first tuple" + "base relation partitioned" + "first node"

**SEARCH**

THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used **receiving first tuple** **base relation partitioned** **first node** **parallel database** **join attribute** **auxiliary relation** **join view**

Found 4 of 171,143

Sort results by relevance  Save results to a Binder  
 Display results expanded form  Search Tips  
 Open results in a new window

Try an [Advanced Search](#)  
 Try this search in [The ACM Guide](#)

Results 1 - 4 of 4

Relevance scale

## 1 [On the development of a site selection optimizer for distributed and parallel database systems](#)



Fotis Barlos, Ophir Frieder  
 December 1993 **Proceedings of the second international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: pdf(1.11 MB)

Additional Information: [full citation](#), [references](#), [index terms](#)

## 2 [Query evaluation techniques for large databases](#)



Goetz Graefe  
 June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

## 3 [Parallel database systems: the future of database processing or a passing fad?](#)



David J. DeWitt, Jim Gray  
 December 1990 **ACM SIGMOD Record**, Volume 19 Issue 4

Publisher: ACM Press

Full text available: pdf(894.07 KB)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The concept of parallel database machines consisting of exotic hardware has been replaced by a fairly conventional shared-nothing hardware base along with a highly parallel dataflow software architecture. Such a design provides speedup and scaleup in processing relational database queries. This paper reviews the techniques used by such systems, and surveys current commercial and research systems.

4 Data placement in shared-nothing parallel database systems □

Manish Mehta, David J. DeWitt

February 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(245.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data placement in shared-nothing database systems has been studied extensively in the past and various placement algorithms have been proposed. However, there is no consensus on the most efficient data placement algorithm and placement is still performed manually by a database administrator with periodic reorganization to correct mistakes. This paper presents the first comprehensive simulation study of data placement issues in a shared-nothing system. The results show that current hardware techn ...

**Keywords:** Declustering, Disk allocation, Resource allocation, Resource scheduling

Results 1 - 4 of 4

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

/7/2006


**PORTAL**  
 USPTO
 
[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)  
 Search:  The ACM Digital Library  The Guide  
 "receiving first tuple" + "base relation partitioned" + "second node" 

THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [receiving first tuple](#) [base relation](#)  
[partitioned](#) [second node](#) [parallel database](#) [join](#)  
[attribute](#) [auxiliary relation](#) [join](#) [view](#)

Found 4 of 175,083

Sort results by relevance  Save results to a Binder  
 Display results expanded form  Search Tips  
 Open results in a new window

Try an [Advanced Search](#)  
 Try this search in [The ACM Guide](#)

Results 1 - 4 of 4

Relevance scale 

1 [On the development of a site selection optimizer for distributed and parallel database systems](#) 

 **Fotis Barlos, Ophir Frieder**  
 December 1993 **Proceedings of the second international conference on Information and knowledge management**

**Publisher:** ACM Press

Full text available:  pdf(1.11 MB)

Additional Information: [full citation](#), [references](#), [index terms](#)

2 [Query evaluation techniques for large databases](#) 

 **Goetz Graefe**  
 June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

**Publisher:** ACM Press

Full text available:  pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

3 [Parallel database systems: the future of database processing or a passing fad?](#) 

 **David J. DeWitt, Jim Gray**  
 December 1990 **ACM SIGMOD Record**, Volume 19 Issue 4

**Publisher:** ACM Press

Full text available:  pdf(894.07 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The concept of parallel database machines consisting of exotic hardware has been replaced by a fairly conventional shared-nothing hardware base along with a highly parallel dataflow software architecture. Such a design provides speedup and scaleup in processing relational database queries. This paper reviews the techniques used by such systems, and surveys current commercial and research systems.

4 Data placement in shared-nothing parallel database systems 

Manish Mehta, David J. DeWitt

February 1997 **The VLDB Journal — The International Journal on Very Large Data**

**Bases**, Volume 6 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(245.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data placement in shared-nothing database systems has been studied extensively in the past and various placement algorithms have been proposed. However, there is no consensus on the most efficient data placement algorithm and placement is still performed manually by a database administrator with periodic reorganization to correct mistakes. This paper presents the first comprehensive simulation study of data placement issues in a shared-nothing system. The results show that current hardware techn ...

**Keywords:** Declustering, Disk allocation, Resource allocation, Resource scheduling

Results 1 - 4 of 4

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

ACM PORTAL  
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search:  The ACM Digital Library  The Guide

"receiving first tuple" + "base relation partitioned" + "first nod"

THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [receiving first tuple](#) [base relation partitioned](#) [first node](#) [parallel database](#) [plural nodes](#) [join attribute](#) [auxiliary relation](#) [join view](#)

Found 1 of 171,143

Sort results by relevance   Save results to a Binder  
 Search Tips  Open results in a new window

Try an [Advanced Search](#)  
Try this search in [The ACM Guide](#)

Results 1 - 1 of 1

Relevance scale

1 [On the development of a site selection optimizer for distributed and parallel database systems](#)

[systems](#)

Fotis Barlos, Ophir Frieder

December 1993 **Proceedings of the second international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: [pdf\(1.11 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

Results 1 - 1 of 1

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

Printable | Help | Log Out


**PORTAL**  
USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search:  The ACM Digital Library  The Guide
"receiving first tuple" + "base relation partitioned" + "first nod
**SEARCH**



 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used **receiving first tuple base relation partitioned first node parallel database plural nodes join attribute auxiliary relation**

Found 4 of 175,083

Sort results by  relevance   
 Display results  expanded form

Save results to a Binder  
 Search Tips  
 Open results in a new window

Try an [Advanced Search](#)  
 Try this search in [The ACM Guide](#)

Results 1 - 4 of 4

Relevance scale

- 1 [On the development of a site selection optimizer for distributed and parallel database systems](#)   
 **systems**  
 Fotis Barlos, Ophir Frieder  
 December 1993 **Proceedings of the second international conference on Information and knowledge management**  
**Publisher:** ACM Press  
 Full text available: [!\[\]\(f7157c5e5a94ed1c3f5c7e79ddb81a5c\_img.jpg\) pdf\(1.11 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

- 2 [Query evaluation techniques for large databases](#)   
 Goetz Graefe  
 June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2  
**Publisher:** ACM Press  
 Full text available: [!\[\]\(bdffe2bdf55c0c3edef3bbaba55e1863\_img.jpg\) pdf\(9.37 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

- 3 [Parallel database systems: the future of database processing or a passing fad?](#)   
 David J. DeWitt, Jim Gray  
 December 1990 **ACM SIGMOD Record**, Volume 19 Issue 4  
**Publisher:** ACM Press  
 Full text available: [!\[\]\(433c20a01a1d31ccc093331a47505c4c\_img.jpg\) pdf\(894.07 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The concept of parallel database machines consisting of exotic hardware has been replaced by a fairly conventional shared-nothing hardware base along with a highly parallel dataflow software architecture. Such a design provides speedup and scaleup in processing relational database queries. This paper reviews the techniques used by such systems, and surveys current commercial and research systems.

4 Data placement in shared-nothing parallel database systems 

Manish Mehta, David J. DeWitt

February 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(245.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data placement in shared-nothing database systems has been studied extensively in the past and various placement algorithms have been proposed. However, there is no consensus on the most efficient data placement algorithm and placement is still performed manually by a database administrator with periodic reorganization to correct mistakes. This paper presents the first comprehensive simulation study of data placement issues in a shared-nothing system. The results show that current hardware techn ...

**Keywords:** Declustering, Disk allocation, Resource allocation, Resource scheduling

Results 1 - 4 of 4

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

4/20/2006

PORTAL
USPTO

Subscribe (Full Service)
Register (Limited Service, Free)
[Login](#)

Search:
 The ACM Digital Library
 The Guide

"receiving first tuple" + "materialized view" + "parallel database"
[SEARCH](#)

## THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used receiving first tuple materialized view parallel database base relation partitioned nodes join attribute auxiliary relation storing join results join view

Found 16 of 171,143

Sort results by relevance  Save results to a Binder  
 Display results expanded form  Search Tips  Open results in a new window

Try an [Advanced Search](#)  
Try this search in [The ACM Guide](#)

Results 1 - 16 of 16

Relevance scale



### 1 [Query evaluation techniques for large databases](#)

Goetz Graefe  
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: [pdf\(9.37 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality



### 2 [Efficient checking of temporal integrity constraints using bounded history encoding](#)

Jan Chomicki  
June 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 2

Publisher: ACM Press

Full text available: [pdf\(2.70 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We present an efficient implementation method for temporal integrity constraints formulated in Past Temporal Logic. Although the constraints can refer to past states of the database, their checking does not require that the entire database history be stored. Instead, every database state is extended with auxiliary relations that contain the historical information necessary for checking constraints. Auxiliary relations can be implemented as materialized relational views.

**Keywords:** active databases, database integrity, integrity constraints, real-time databases, temporal databases, temporal logic, triggers

**3 The complexity of acyclic conjunctive queries** □

 May 2001 **Journal of the ACM (JACM)**, Volume 48 Issue 3

Publisher: ACM Press

Full text available:  pdf(566.16 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

**Keywords:** CSP, LOGCFL, acyclic hypergraph, algorithm, bounded treewidth, conjunctive query, constraint, constraint satisfaction problem, database theory, degree of cyclicity, hinge, join tree, parallel algorithm, query containment, query-depth, subsumption, tree query

**4 Research papers: streams and pipelined processing: QPipe: a simultaneously pipelined relational query engine** □

 Stavros Harizopoulos, Vladislav Shkapenyuk, Anastassia Ailamaki

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available:  pdf(506.36 KB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Relational DBMS typically execute concurrent queries independently by invoking a set of operator instances for each query. To exploit common data retrievals and computation in concurrent queries, researchers have proposed a wealth of techniques, ranging from buffering disk pages to constructing materialized views and optimizing multiple queries. The ideas proposed, however, are inherently limited by the query-centric philosophy of modern engine designs. Ideally, the query engine should proactive ...

**5 Physical design: Experimental evidence on partitioning in parallel data warehouses** □

 Pedro Furtado

November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP**

Publisher: ACM Press

Full text available:  pdf(260.86 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Parallelism can be used for major performance improvement in large Data warehouses (DW) with performance and scalability challenges. A simple low-cost shared-nothing architecture with horizontally fully-partitioned facts can be used to speedup response time of the data warehouse significantly. However, extra overheads related to processing large replicated relations and repartitioning requirements between nodes can significantly degrade speedup performance for many query patterns if special c ...

**Keywords:** data warehouse, parallel, performance

**6 Compensation-based on-line query processing** □

 V. Srinivasan, Michael J. Carey

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

It is well known that using conventional concurrency control techniques for obtaining serializable answers to long-running queries leads to an unacceptable drop in system

performance. As a result, most current DBMSs execute such queries under a reduced degree of consistency, thus providing non-serializable answers. In this paper, we present a new and highly concurrent approach for processing large decision support queries in relational databases. In this new approach, called compensation-ba ...

7 PODS invited talk: Models and issues in data stream systems

 Brian Babcock, Shivnath Babu, Mayur Datar, Rajeev Motwani, Jennifer Widom  
June 2002 **Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  pdf(257.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this overview paper we motivate the need for and research issues arising from a new model of data processing. In this model, data does not take the form of persistent relations, but rather arrives in multiple, continuous, rapid, time-varying *data streams*. In addition to reviewing past work relevant to data stream systems and current projects in the area, the paper explores topics in stream query languages, new requirements and challenges in query processing, and algorithmic issues.

8 Perspectives on database theory

 Mihalis Yannakakis  
September 1996 **ACM SIGACT News**, Volume 27 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.13 MB) Additional Information: [full citation](#), [index terms](#)

9 Paper session DB-9 (databases): query processing 1: Selectivity-based partitioning:

 a divide-and-union paradigm for effective query optimization

Neoklis Polyzotis

October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**

Publisher: ACM Press

Full text available:  pdf(253.90 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Modern query optimizers select an efficient join ordering for a physical execution plan based essentially on the average join selectivity factors among the referenced tables. In this paper, we argue that this "monolithic" approach can miss important opportunities for the effective optimization of relational queries. We propose *selectivity-based partitioning*, a novel optimization paradigm that takes into account the join correlations among *relation fragments* in order to essen ...

10 Industrial sessions: big data: Automating physical database design in a parallel database

 Jun Rao, Chun Zhang, Nimrod Megiddo, Guy Lohman  
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(1.38 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Physical database design is important for query performance in a shared-nothing parallel database system, in which data is horizontally partitioned among multiple independent nodes. We seek to automate the process of data partitioning. Given a workload of SQL statements, we seek to determine automatically how to partition the base data across multiple nodes to achieve overall optimal (or close to optimal) performance for that

workload. Previous attempts use heuristic rules to make those decision ...

11 Query reformulation using materialized views in data warehouse environment

 Jae-young Chang, Sang-goo Lee  
November 1998 **Proceedings of the 1st ACM international workshop on Data warehousing and OLAP**

Publisher: ACM Press

Full text available:  pdf(763.85 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 Parallel database systems: the future of database processing or a passing fad?

 David J. DeWitt, Jim Gray  
December 1990 **ACM SIGMOD Record**, Volume 19 Issue 4

Publisher: ACM Press

Full text available:  pdf(894.07 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The concept of parallel database machines consisting of exotic hardware has been replaced by a fairly conventional shared-nothing hardware base along with a highly parallel dataflow software architecture. Such a design provides speedup and scaleup in processing relational database queries. This paper reviews the techniques used by such systems, and surveys current commercial and research systems.

13 Materialized view maintenance and integrity constraint checking: trading space for time

 Kenneth A. Ross, Divesh Srivastava, S. Sudarshan  
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.34 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We investigate the problem of incremental maintenance of an SQL view in the face of database updates, and show that it is possible to reduce the total time cost of view maintenance by materializing (and maintaining) additional views. We formulate the problem of determining the optimal set of additional views to materialize as an optimization problem over the space of possible view sets (which includes the empty set). The optimization problem is harder than query optimization since it has to deal ...

14 Adapting materialized views after redefinitions

 Ashish Gupta, Inderpal S. Mumick, Kenneth A. Ross  
May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data SIGMOD '95**, Volume 24 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.24 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We consider a variant of the view maintenance problem: How does one keep a materialized view up-to-date when the view definition itself changes? Can one do better than recomputing the view from the base relations? Traditional view maintenance tries to maintain the materialized view in response to modifications to the base relations; we try to "adapt" the view in response to changes in the view definition. Such techniques are needed for applications where the user can change queries dynamically an ...

15 Data placement in shared-nothing parallel database systems

Manish Mehta, David J. DeWitt  
February 1997 **The VLDB Journal — The International Journal on Very Large Data**

**Bases**, Volume 6 Issue 1

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  [pdf\(245.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data placement in shared-nothing database systems has been studied extensively in the past and various placement algorithms have been proposed. However, there is no consensus on the most efficient data placement algorithm and placement is still performed manually by a database administrator with periodic reorganization to correct mistakes. This paper presents the first comprehensive simulation study of data placement issues in a shared-nothing system. The results show that current hardware techn ...

**Keywords:** Declustering, Disk allocation, Resource allocation, Resource scheduling

- 16 [Research sessions: indexing and tuning: Integrating vertical and horizontal partitioning into automated physical database design](#) 

 Sanjay Agrawal, Vivek Narasayya, Beverly Yang  
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

**Publisher:** ACM Press

Full text available:  [pdf\(181.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

In addition to indexes and materialized views, horizontal and vertical partitioning are important aspects of physical design in a relational database system that significantly impact performance. Horizontal partitioning also provides manageability; database administrators often require indexes and their underlying tables partitioned identically so as to make common operations such as backup/restore easier. While partitioning is important, incorporating partitioning makes the problem of automatin ...

Results 1 - 16 of 16

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

11/20/06 4/20/06


**PORTAL**  
 USPTO
 
[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)  
**Search:**  The ACM Digital Library  The Guide  
 "receiving first tuple" + "materialized view" + "parallel database" 

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [receiving first tuple](#) [materialized view](#) [parallel database](#) [base relation](#) [partitioned nodes](#) [join attribute](#) [auxiliary relation](#) [storing join view](#)

Found 16 of 171,143

Sort results by relevance  Save results to a Binder  
 Display results expanded form  Search Tips  Open results in a new window

Try an [Advanced Search](#)  
 Try this search in [The ACM Guide](#)

Results 1 - 16 of 16

Relevance scale

**1** [Query evaluation techniques for large databases](#)

Goetz Graefe  
 June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

**Publisher:** ACM PressFull text available: [pdf\(9.37 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

**2** [Efficient checking of temporal integrity constraints using bounded history encoding](#)

Jan Chomicki  
 June 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 2

**Publisher:** ACM PressFull text available: [pdf\(2.70 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We present an efficient implementation method for temporal integrity constraints formulated in Past Temporal Logic. Although the constraints can refer to past states of the database, their checking does not require that the entire database history be stored. Instead, every database state is extended with auxiliary relations that contain the historical information necessary for checking constraints. Auxiliary relations can be implemented as materialized relational views.

**Keywords:** active databases, database integrity, integrity constraints, real-time databases, temporal databases, temporal logic, triggers

**3 The complexity of acyclic conjunctive queries**

 May 2001 **Journal of the ACM (JACM)**, Volume 48 Issue 3

**Publisher:** ACM Press

Full text available:  pdf(566.16 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

**Keywords:** CSP, LOGCFL, acyclic hypergraph, algorithm, bounded treewidth, conjunctive query, constraint, constraint satisfaction problem, database theory, degree of cyclicity, hinge, join tree, parallel algorithm, query containment, query-depth, subsumption, tree query

**4 Research papers: streams and pipelined processing: QPipe: a simultaneously pipelined relational query engine**

 Stavros Harizopoulos, Vladislav Shkapenyuk, Anastassia Ailamaki

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

**Publisher:** ACM Press

Full text available:  pdf(506.36 KB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Relational DBMS typically execute concurrent queries independently by invoking a set of operator instances for each query. To exploit common data retrievals and computation in concurrent queries, researchers have proposed a wealth of techniques, ranging from buffering disk pages to constructing materialized views and optimizing multiple queries. The ideas proposed, however, are inherently limited by the query-centric philosophy of modern engine designs. Ideally, the query engine should proactive ...

**5 Physical design: Experimental evidence on partitioning in parallel data warehouses**

 Pedro Furtado

November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP**

**Publisher:** ACM Press

Full text available:  pdf(260.86 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Parallelism can be used for major performance improvement in large Data warehouses (DW) with performance and scalability challenges. A simple low-cost shared-nothing architecture with horizontally fully-partitioned facts can be used to speedup response time of the data warehouse significantly. However, extra overheads related to processing large replicated relations and repartitioning requirements between nodes can significantly degrade speedup performance for many query patterns if special c ...

**Keywords:** data warehouse, parallel, performance

**6 Compensation-based on-line query processing**

 V. Srinivasan, Michael J. Carey

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

**Publisher:** ACM Press

Full text available:  pdf(1.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

It is well known that using conventional concurrency control techniques for obtaining serializable answers to long-running queries leads to an unacceptable drop in system

performance. As a result, most current DBMSs execute such queries under a reduced degree of consistency, thus providing non-serializable answers. In this paper, we present a new and highly concurrent approach for processing large decision support queries in relational databases. In this new approach, called compensation-based ...

7 PODS invited talk: Models and issues in data stream systems

 Brian Babcock, Shivnath Babu, Mayur Datar, Rajeev Motwani, Jennifer Widom  
June 2002 **Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  pdf(257.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this overview paper we motivate the need for and research issues arising from a new model of data processing. In this model, data does not take the form of persistent relations, but rather arrives in multiple, continuous, rapid, time-varying *data streams*. In addition to reviewing past work relevant to data stream systems and current projects in the area, the paper explores topics in stream query languages, new requirements and challenges in query processing, and algorithmic issues.

8 Perspectives on database theory

 Mihalis Yannakakis  
September 1996 **ACM SIGACT News**, Volume 27 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.13 MB) Additional Information: [full citation](#), [index terms](#)

9 Paper session DB-9 (databases): query processing 1: Selectivity-based partitioning:

 a divide-and-union paradigm for effective query optimization

Neoklis Polyzotis

October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**

Publisher: ACM Press

Full text available:  pdf(253.90 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Modern query optimizers select an efficient join ordering for a physical execution plan based essentially on the average join selectivity factors among the referenced tables. In this paper, we argue that this "monolithic" approach can miss important opportunities for the effective optimization of relational queries. We propose *selectivity-based partitioning*, a novel optimization paradigm that takes into account the join correlations among *relation fragments* in order to ...

10 Industrial sessions: big data: Automating physical database design in a parallel database

 Jun Rao, Chun Zhang, Nimrod Megiddo, Guy Lohman  
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(1.38 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Physical database design is important for query performance in a shared-nothing parallel database system, in which data is horizontally partitioned among multiple independent nodes. We seek to automate the process of data partitioning. Given a workload of SQL statements, we seek to determine automatically how to partition the base data across multiple nodes to achieve overall optimal (or close to optimal) performance for that

workload. Previous attempts use heuristic rules to make those decision ...

11 Query reformulation using materialized views in data warehouse environment

 Jae-young Chang, Sang-goo Lee  
November 1998 **Proceedings of the 1st ACM international workshop on Data warehousing and OLAP**

Publisher: ACM Press

Full text available:  pdf(763.85 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 Parallel database systems: the future of database processing or a passing fad?

 David J. DeWitt, Jim Gray  
December 1990 **ACM SIGMOD Record**, Volume 19 Issue 4

Publisher: ACM Press

Full text available:  pdf(894.07 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The concept of parallel database machines consisting of exotic hardware has been replaced by a fairly conventional shared-nothing hardware base along with a highly parallel dataflow software architecture. Such a design provides speedup and scaleup in processing relational database queries. This paper reviews the techniques used by such systems, and surveys current commercial and research systems.

13 Materialized view maintenance and integrity constraint checking: trading space for

 time  
Kenneth A. Ross, Divesh Srivastava, S. Sudarshan  
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.34 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We investigate the problem of incremental maintenance of an SQL view in the face of database updates, and show that it is possible to reduce the total time cost of view maintenance by materializing (and maintaining) additional views. We formulate the problem of determining the optimal set of additional views to materialize as an optimization problem over the space of possible view sets (which includes the empty set). The optimization problem is harder than query optimization since it has to deal ...

14 Adapting materialized views after redefinitions

 Ashish Gupta, Inderpal S. Mumick, Kenneth A. Ross  
May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data SIGMOD '95**, Volume 24 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.24 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We consider a variant of the view maintenance problem: How does one keep a materialized view up-to-date when the view definition itself changes? Can one do better than recomputing the view from the base relations? Traditional view maintenance tries to maintain the materialized view in response to modifications to the base relations; we try to "adapt" the view in response to changes in the view definition. Such techniques are needed for applications where the user can change queries dynamically an ...

15 Data placement in shared-nothing parallel database systems

Manish Mehta, David J. DeWitt  
February 1997 **The VLDB Journal — The International Journal on Very Large Data**

**Bases**, Volume 6 Issue 1

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  [pdf\(245.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data placement in shared-nothing database systems has been studied extensively in the past and various placement algorithms have been proposed. However, there is no consensus on the most efficient data placement algorithm and placement is still performed manually by a database administrator with periodic reorganization to correct mistakes. This paper presents the first comprehensive simulation study of data placement issues in a shared-nothing system. The results show that current hardware techn ...

**Keywords:** Declustering, Disk allocation, Resource allocation, Resource scheduling

16 [Research sessions: indexing and tuning: Integrating vertical and horizontal](#)



[partitioning into automated physical database design](#)

 Sanjay Agrawal, Vivek Narasayya, Beverly Yang

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

**Publisher:** ACM Press

Full text available:  [pdf\(181.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

In addition to indexes and materialized views, horizontal and vertical partitioning are important aspects of physical design in a relational database system that significantly impact performance. Horizontal partitioning also provides manageability; database administrators often require indexes and their underlying tables partitioned identically so as to make common operations such as backup/restore easier. While partitioning is important, incorporating partitioning makes the problem of automatin ...

Results 1 - 16 of 16

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)